



Europäisches Patentamt
European Patent Office
Office européen des brevets

⑯ Publication number:

0 262 938
A1

⑰

EUROPEAN PATENT APPLICATION

㉑ Application number: 87308638.3

㉓ Int. Cl. 4: G 06 F 15/38
G 06 F 3/16

㉒ Date of filing: 29.09.87

㉔ Priority: 03.10.86 GB 8623839
01.05.87 GB 8710376

㉕ Date of publication of application:
06.04.88 Bulletin 88/14

㉖ Designated Contracting States:
AT BE CH DÉ ES FR GB GR IT LI LU NL SE

㉗ Applicant: BRITISH TELECOMMUNICATIONS public limited company
British Telecom Centre, 81 Newgate Street
London EC1A 7AJ (GB)

㉘ Inventor: Stentiford, Frederick Warwick Michael
Sheepstor Boyton
Woodbridge Suffolk IP12 3LH (GB)

Steer, Martin George
Clover Cottage
Offton Ipswich Suffolk IP8 4SE (GB)

㉙ Representative: Roberts, Simon Christopher et al
BRITISH TELECOM Intellectual Property Unit 13th Floor
151, Gower Street
London, WC1E 6BA (GB)

㉚ Language translation system.

㉛ A language translation system for translating phrases from a first language into a second language comprises a store holding a collection of phrases in the second language. Phrases input in the first language are each characterised on the basis of one or more keywords, and the corresponding phrase in the second language is output. Such a phrasebook approach enables what is effectively rapid and accurate translation, even from speech. Since the phrases in the second language are prepared in advance and held in store, there need be no problems of poor translation or ungrammatical construction. The output may be in text, or, using speech synthesis, in voiced form.

With appropriate choice of keywords it is possible to characterise a large number of relatively long and complex phrases with just a few keywords.

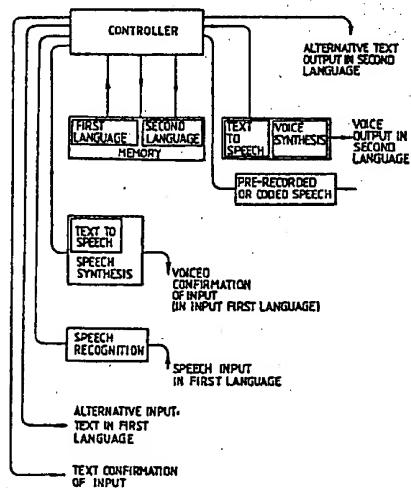


Fig.1

EP 0 262 938 A1

Description**LANGUAGE TRANSLATION SYSTEM**

This invention relates to a system for translating phrases from a first language to a second language, and in particular but not exclusively to such a system for producing speech in a second language from speech in a first language.

A machine which can rapidly and automatically translate languages, particularly speech, has been sought for many years. However, even with the tremendous progress in computing, speech recognition and speech synthesis in recent years, such machines remain the stuff of dreams and fiction.

Considerable research has been carried out on computer systems for the automatic translation of text. Apart from a few very restricted applications (eg translation of weather forecasts), no product exists which can automatically produce accurate translations and hence replace human translators. The problems of translating speech are compounded by the errors of speech recognition, the additional information intonation, stress etc and the inexactness of speech itself.

Unfortunately, existing text language translation packages are all deficient in some way or another and do not meet the requirements of a system translating speech-to-speech. Most such packages have been designed as an aid for professional translators, and produce outputs which have to be post-edited before being presentable in their target language. Most packages are either menu-driven and interactive or operate in a slow batch processing mode, neither of which is suitable for "real-time" speech operation. Translation packages also tend to be unreliable, as idioms and other exceptions can easily cause erroneous output: the user has no guarantee that the output is correctly translated. Existing systems are also very CPU intensive, making them expensive to run and hence unsuitable for many cost sensitive applications.

The present invention seeks to provide a translation system in which these deficiencies and disadvantages are mitigated.

According to the present invention there is provided a system for translating phrases from a first language into a second language, comprising: a store holding a collection of phrases in the second language; input means to accept a phrase in the first language; output means to output in the second language a phrase comprising one from said collection of phrases; characterisation means to determine which of said collection of phrases corresponds to said input phrase; means responsive to said characterisation means to control the output means and to ensure the outputting of the phrase from said collection which corresponds to said input phrase.

Such a system provides very quick translation, the time required being that to identify/characterise the input phrase and that to look up the 'answer' in the second language.

The system can also be implemented to give the user providing the input confirmation that she/he

has been recognised/understood correctly by the system, which is of course particularly important in speech translation systems.

Once it has been confirmed to the user that his message has been correctly characterised, accuracy of translation is ensured because the stored collection of phrases consists only of previously made accurate translations.

The system also makes possible rapid translation into several second languages simultaneously; essentially all that need be added are further stores holding collections of phrases in each of the additional second languages.

Embodiments of the invention will now be described with reference to the accompanying drawings in which:

Figure 1 is a block diagram showing the principal components of a system according to the invention.

The present invention is based on our appreciation that it is possible to characterise and capture the semantic content of a large number of distinct phrases by means of a very much smaller number of keywords. With appropriate selection of the keywords it is possible to use existing, commercially available speech recognisers, which are only capable of recognising considerably fewer words than would be contained in a usefully large set of phrases, to characterise and differentiate a large set of phrases.

The performance of the translation system as a whole therefore rests on the ability of those keywords to correctly distinguish between phrases. The greater the separation of phrases achieved, the greater the system's tolerance to recognition errors, and also discrepancies introduced by the speaker himself.

The Selection of Keywords

A suitable search procedure is as follows:

1. Order each of the K words in the N phrases of interest according to the word's frequency of occurrence in the phrases.
2. Select the M most frequently occurring words as the initial keyword list, where M is the number of words in the vocabulary of the speech recogniser.
3. The presence or otherwise of each keyword in each phrase is then determined. The number of phrases (E) which are not distinguished by the keywords are counted.
4. Let i = 1.
5. A keyword is temporarily deleted from the list and the new value (E') of E is computed.
6. The score E'-E is assigned to the keyword which was temporarily deleted; this being a measure of worsening of the performance after the renewal of the keyword, and hence its contribution to the overall performance. [In effect, this measure is used to ensure that each keyword contributes to the separation of as many phrase pairs as possible but without

simply duplicating the function of others.]

7. Temporarily deleted keywords are replaced and the process is repeated for each of the M keywords.

8. The word with the lowest score is removed from the current keyword list.

9. The M + ith most frequent word is then used to replace the removed word, and then a new E is calculated.

10. If the new E indicates an improved performance over the previous E then i is incremented and the process repeated from step 5 unless $M + i > K$ in which case the process stops. Otherwise the M + ith word is rejected; i is incremented and the process repeated from step 9 unless $M + i > K$ in which case the word last removed in step 8 is replaced and the process stops.

The final keyword list contains the optimal set of M single keywords for phrase identification.

Further iterations starting with the best M words from the previous iteration may yield further improvements in phrase separation. Heuristics other than frequency ordering may be used to provide the succession of candidate words in step 1, especially if a priori linguistic information is available. In addition, it is likely that the words towards the bottom of the occurrence list will not appreciably aid separation of phrases, and it may therefore not be worth searching through more than say the upper third or upper half of the occurrence list.

It is sometimes the case that most phrases are distinguished and E becomes very close to zero quite early in the search. Further improvements are obtained in these cases by computing E on the basis that phrases are only considered distinguished if more than one keyword is different. This ensures that most phrases are separated by more than a minimum number of keywords and provides some immunity to speech recognition errors.

During the search it becomes clear that several classes of phrase are never going to be separated unless the keyword vocabulary is extended. These "clusters" or groups of phrases tend to differ only by a single word or subordinate string of word (eg dates in business letters), and are candidates derived automatically for use in the preparation of keyword subvocabularies (detailed below).

It is apparent that the recognition of single keywords takes no account of word order and the additional meaning that it may contain. The presence or otherwise of key pairs (or other multiples) of words with various separations between them can therefore also be used to improve the effectiveness of the single keyword set. This has the advantage in speech recognition that the performance may be improved without increasing the recognition vocabulary. In a text application further improvements can be obtained by generalising the keywords to include punctuation, parts of words, and combinations of words and parts of words, eg "-ing * bed" (where * can be any word) would be present in "making the bed" and "selling a bed".

The use of pairs of keywords (eg we * * to) enhances the value of the component single words if

further phrases confusions are resolved. The search for word pairs which are not necessarily contiguous but separated by different numbers of other words, again begins with the preparation of a frequency ordering. Word pairs with both component words in the M keywords are made from the ordered list if they resolve any remaining phrase confusions. The final list of single keywords and pairs of keywords are each scored as before and an overall phrase confusion score E computed.

The search now begins for better performing word pairs where one or both of the component keywords are not in the current keyword list. The next word pair candidate is taken from the top of the frequency ordering and appended to the keyword list. The single keywords in the appended word pair which are not already present are also added and an equal number of the worst performing single keywords deleted. This may cause other word pairs to be deleted if their component words are no longer present. A new value (E') of E is computed. If an improvement is obtained and $E' < E$, the most recent modifications of the keyword list are retained, otherwise the list is restored to its previous state. Further word pairs are processed from the frequency ordering, although as with the single keyword search, other heuristics may be used to provide candidate word pairs.

It is worth observing that some keywords contribute more to the overall performance through their participation in several word groups than by themselves.

The method extends to larger keyword groupings (> 2 words), but as the frequency of occurrence decreases, the contribution to the resolution of phrase confusions are only significant in a very large corpus of phrases.

The quantity of computation involved in the search for keywords increases with the number of keywords and the number of phrases. This may be reduced by first running the algorithm on a subset of phrases which are confused or very close to being confused. The keywords and their scores so obtained provide a more efficient ordering of candidate keywords to the main algorithm which will work with a more complete set of phrases.

In a speech recognition application some words which are not in the keyword set can generate many spurious keyword recognitions. e.g. occurrences of the word "I" may be always recognised as the keyword "by". If however, the groups of confused words are considered as synonymous before the search for keywords begins and in the subsequent phrase identification, the actual phrase separations should not be affected by this problem. Furthermore because the frequency of such synonymous words taken together is necessarily higher than that of the separate words, a greater quantity of phrasal information is normally associated with their detection.

The use of keywords may be extended to keyword-parts (e.g. phonemes) which occur again with higher frequency and which bear more phrase distinguishing information than the whole words. Moreover the identification of certain word-parts in

continuous speech is often easier than complete words, and is therefore preferable in a translation system which accepts continuous speech input. Throughout this specification the word "keyword" is for the sake of brevity used to refer to both whole keywords and to parts of keywords.

Many classes of phrase only differ from each other in subordinate phrases and clauses which may contain details of dates, time, prices, item, names or other groups of words. It may be that the vocabulary of a speech recogniser is sufficient to assign a phrase to a particular class or group of phrases but is not large enough to hold sufficient keywords to separate the subordinate structures. Furthermore it is quite possible that the total vocabulary required to separate the phrase classes and the subordinate structure contains many more words which are easily confused. This means that even if the capacity of the recogniser was sufficient to cover the whole vocabulary, the performance would be too low to obtain reliable phrase and subordinate phrase identification. It is an advantage of the method according to the invention that the original utterance or some transform of the original utterance may be stored in a buffer and the recognition process may be repeated, once the phrase class has been determined, using the set of keywords which are expected in the subordinate word strings particular to that phrase class. In this way the recognition apparatus never has to cope with the total vocabulary, with its many potential word confusions, at once, but appears to the user to do so. It should be noted that the speed of the second recognition process is not limited by the speed of the original utterance and can in principle be carried out much faster than real time and hence not necessarily introduce noticeable delays. The iterations of recognition may be carried out as many times as is necessary to identify the required phrase and its substructure. It thus becomes possible to 'nest' the recognition process, the phrase being characterised in numerous separate stages, the recogniser at each stage drawing on a different vocabulary of keywords.

Many, although not all, subordinate word strings will be context independent in the source language. This is because positions for subordinate word strings are only designated as such if several alternatives are possible making tight contextual dependence less likely for any one of them. In addition contextual importance would imply that there were dependencies between words which were inside and outside the potential subordinate string and hence there would be scope for keywords to distinguish the whole phrase without the use of words inside the string. This is illustrated in phrases containing changing dates in which there is rarely any word change necessary in the phrase apart from the date itself. (It is for future research to demonstrate the conjecture that such context independence is generally invariant between languages and use it to extend phrasebook translation indefinitely.)

This particular aspect of the invention also has significant benefits when employed for the translation of text where the computational costs of searching large dictionaries can be reduced dramati-

cally by using a similar hierarchy of smaller dictionaries and phrasebooks. Some subordinate phrases do not need to be translated and often in these cases it would not in general be possible to recognise automatically the words in these phrases. The commonest case of this occurs in utterances which make reference to labels such as proper nouns: eg. "Can I speak to Mr Smith please?". As before, the system can identify the phrase class together with the locations of words in the buffer which correspond to the label reference. The processing of such label reference words during translation is then simply the transmission of the original acoustic signal in the appropriate place in the target language utterance. Clearly it is desirable that the synthesised target language voice should match the voice of the original speaker and it is a requirement of text-to-speech synthesisers that certain speech parameters can be set so that such matching can be achieved as far as possible (eg. old/young, male/female).

So that the user can be sure that the correct phrase will be output in the target language, the system indicates what phrase in the input language it will translate. In order to be able to do this, the system is provided with a store holding the full repertoire of phrases in the input language.

Preferably in the system the phrases are stored as text, in for example ASCII coded form, since that reduces the storage requirement very considerably compared to that needed for conventionally companded or non-companded speech. Where speech output is required, the text is retrieved from store and passed to a text to speech converter and speech synthesiser. With ASCII coded text storage, 1 byte per character is needed, which means that about 10,000 phrases could be stored with half a megabyte of storage. Hence a system providing translation of about 10,000 phrases would require about 1 megabyte of storage -which is easily provided on hard disc. Two-way communication is possible using two symmetrically constructed translation systems. This has the advantage that each unit is only concerned with recognising and synthesising words in the language of the person operating that unit. Communication with the second unit is by means of a protocol which specifies the phrase and the contents of any subordinate phrases. The protocol is independent of language and hence allows messages to be transmitted without the need to identify the target language. In addition it allows people using many different languages to receive simultaneously translations from the output of a single unit.

EXAMPLE

A demonstration system, connected to a telephone network, has been run to demonstrate the feasibility of the phrase-book approach. The demonstration system uses a Votan speech recogniser, an Infovox speech synthesiser and an IBM PC XT computer.

The Votan speech recogniser is capable of recognizing up to 64 continuously spoken words over a telephone network. Allowing for system

control words such as "yes", "no", "quit" and "enter", upto 60 words can be chosen to be keywords. None of the system control words are allowed to appear in the input phrases, so where it is possible it may be preferable to use control buttons or keys rather than spoken commands.

The store of phrases consists of 400 English phrases and their French equivalents.

The English phrases contain around 1100 different words. To put these numbers in context, a standard phrasebook of business expressions would typically contain this number of phrases.

After running keyword extraction software based on the principles outlined above, 60 keywords were chosen which successfully separated all the phrases. Of the 400 phrases, only 32 were distinguished by just a single word (those 32 phrases being in 16 pairs).

The demonstration system, on recognising the keyword(s), accesses the appropriate phrase, confirms it (orally) with the user and outputs the French equivalent via a text to speech synthesiser.

It is important to note that text-to-speech synthesis is not essential to this invention. It is quite feasible, indeed advantageous, to synthesise target language speech from pre-recorded or coded words and phrases. This has the advantage that such speech may be recorded by the user and hence will acoustically match any embedded speech, and removes the need for text-to-speech synthesis. The approach also removes the need for text-to-speech synthesis in the languages of important countries where such technology is unlikely to produce useable hardware in the immediate future -for example Hindi and Arabic.

In addition to speech-to-speech translation, the present invention is of course applicable to text-to-text, text-to-speech or speech-to-text translation. A particularly useful application is in the field of office automation, where a speech activated foreign language text producing machine could readily be implemented. Essentially, such a machine would use the speech recogniser, software and control system described above, but output the 2nd language text to a printer or telex or other telecommunications link. It would of course be a simple matter to provide the standard phrases of everyday business correspondence in several languages.

Claims

1 A system for translating phrases from a first language into a second language, comprising: a store holding a collection of phrases in the second language; input means to accept a phrase in the first language; output means to output in the second language a phrase comprising one from said collection of phrases; characterisation means to determine which of said collection of phrases corresponds to said input phrase; and means responsive to said characterisation means to control the output means and to ensure the outputting of the

phrase from said collection which corresponds to said input phrase.

2 A system as claimed in claim 1, the system comprising first and second terminals operably connected via a data link, the first terminal comprising said input means and said characterisation means; the second terminal comprising said store and said output means; wherein said first terminal further comprises means to generate a message for transmission to said second terminal via said data link, which message indicates which of said collection of phrases corresponds to said input phrase.

3 A system as claimed in claim 1 or claim 2, in which the characterisation means determines the correspondence of phrases on the basis of the presence in the input phrase of a keyword, keywords, or keyword-part; the characterisation means applying a first vocabulary to determine to which phrase or group of phrases, if any, from said collection of phrases the input phrase corresponds; wherein, in the cases that the input phrase is found to correspond to an undetermined one of a group of phrases, the characterisation means applies a second vocabulary to determine to which of a group of phrases the input phrase corresponds.

4 A system as claimed in anyone of the preceding claims wherein the characterisation means comprises a speech recogniser.

5 A system as claimed in any one of the preceding claims, wherein said input means is capable of accepting spoken inputs, and said output means provides voiced outputs.

6 A system as claimed in anyone of the preceding claims wherein means are provided to enable portions of said input phrase to be passed untranslated to said output means for outputting as part of the phrase in the second language.

7 Apparatus for implementing a system as claimed in any one of the preceding claims.

8 Apparatus as claimed in claim 7 further comprising a keyboard for providing an input message to said input means, and means to provide a text output in said second language.

9 Apparatus as claimed in claim 7 or claim 8 for providing translations from a first language into any one of plurality of second languages, a further collection of phrases being provided in a store in each of said plurality of second languages.

0262938

U.S. Patent and
Trademark Office
Newly filed / Nouvellement déposé

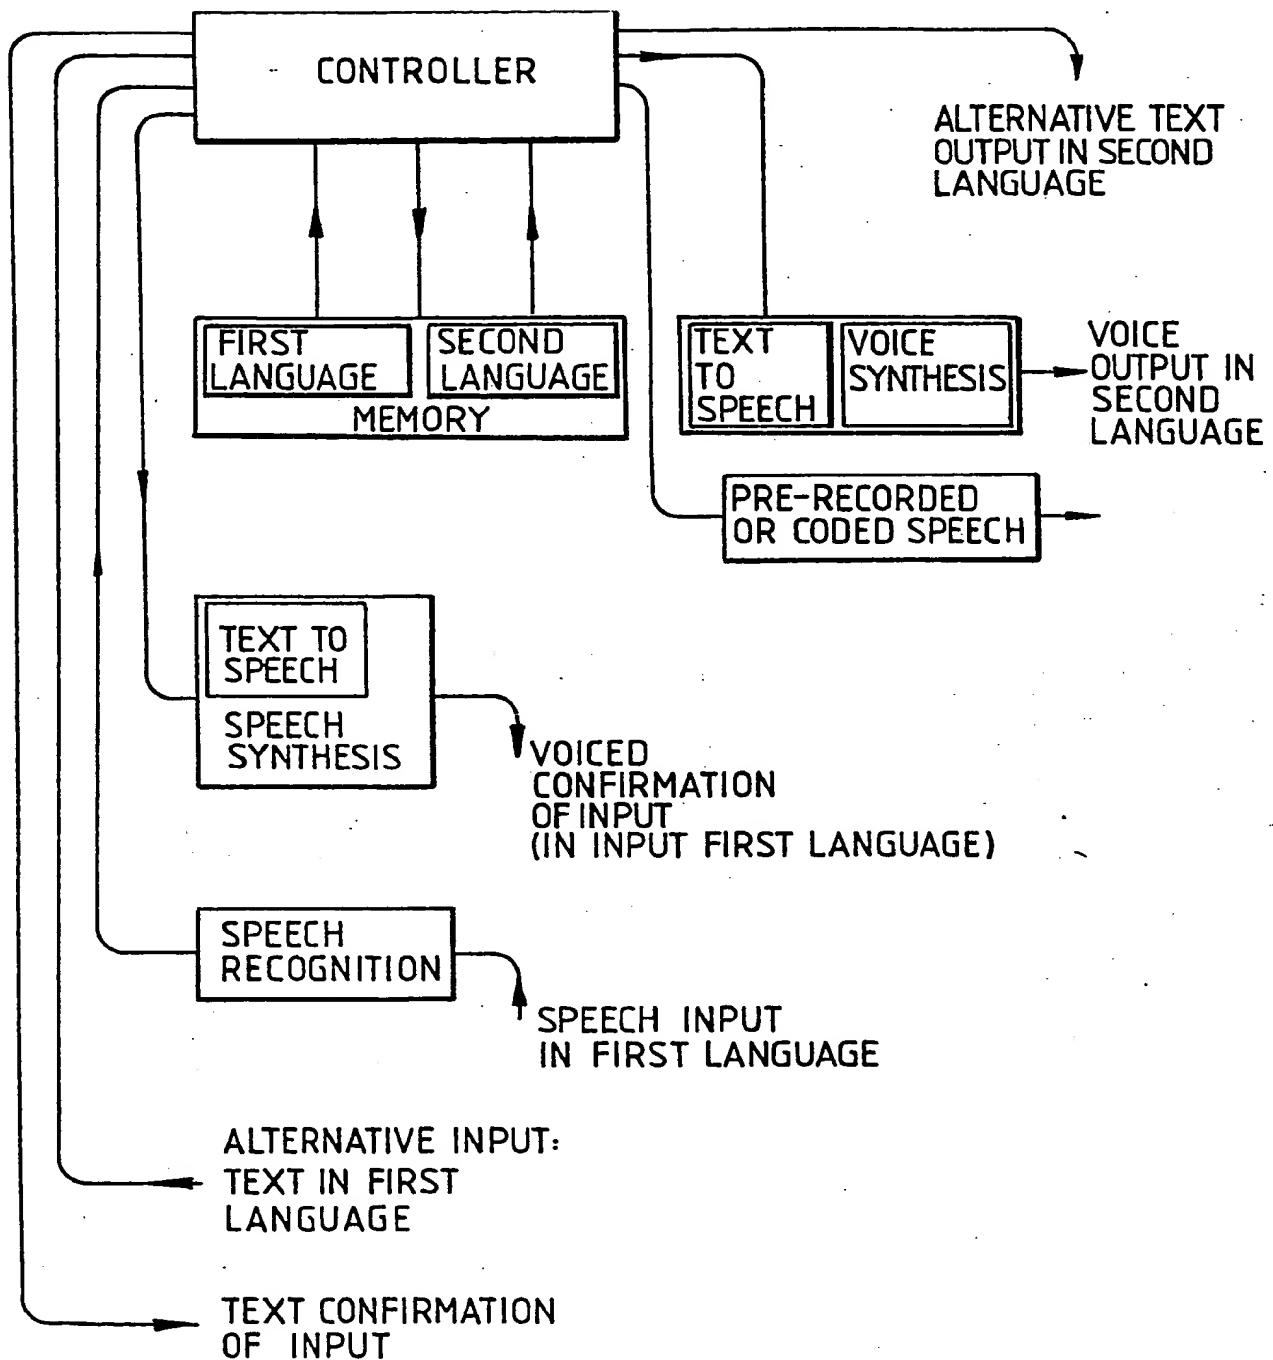


Fig.1